

---

# REDUCING THE COMPUTATIONAL BURDEN OF DEEP LEARNING WITH RECURSIVE LOCAL REPRESENTATION ALIGNMENT

---

**Alexander Ororbia\***  
Rochester Institute of Technology  
Rochester, NY 14623  
ago@cs.rit.edu

**Ankur Mali\***  
The Pennsylvania State University  
State College, PA 16801  
aam35@psu.edu

**Daniel Kifer**  
The Pennsylvania State University  
State College, PA 16801  
duk17@psu.edu

**C. Lee Giles**  
The Pennsylvania State University  
State College, PA 16801  
c1g20@psu.edu

## ABSTRACT

Training deep neural networks on large-scale datasets requires significant hardware resources whose costs (even on cloud platforms) put them out of reach of smaller organizations, groups, and individuals. Backpropagation (backprop), the workhorse for training these networks, is an inherently sequential process that is difficult to parallelize. Furthermore, it requires researchers to continually develop various tricks, such as specialized weight initializations and activation functions, in order to ensure a stable parameter optimization. Our goal is to seek an effective, parallelizable alternative to backprop that can be used to train deep networks. In this paper, we propose a gradient-free learning procedure, *recursive local representation alignment*, for training large-scale neural architectures. Experiments with deep residual networks on CIFAR-10 and the massive-scale benchmark, ImageNet, show that our algorithm generalizes as well as backprop while converging sooner due to weight updates that are parallelizable and computationally less demanding. This is empirical evidence that a backprop-free algorithm can scale up to larger datasets. Another contribution is that we also significantly reduce total parameter count of our networks by utilizing fast, fixed noise maps in place of convolutional operations without compromising generalization.

## 1 Introduction

At the heart of training artificial neural networks (ANNs) is the calculation of adjustments that need to be made to parameters given some data. This calculation is used in tandem with an optimization procedure, such as a stochastic hill climbing procedure, to then alter the ANN's actual parameters in order to ensure it makes better future predictions. This adjustment process entails using an algorithm that can conduct credit assignment, i.e., the task of determining the contribution that individual neuronal units (within the ANN) make to the system's overall error. To conduct credit assignment and compute weight updates in state-of-the-art networks today, back-propagation of errors (backprop) [31] is the popular algorithm of choice. While backprop provides a theoretical basis for training networks, i.e. gradient descent, it also presents practical challenges, particularly due to the problems of exploding/vanishing gradients [9].

In order to deal with the problems posed by backprop, researchers must resort to tricks and heuristics, e.g., careful initialization of weights, often following from a network-specific analysis of backprop's learning dynamics [9, 11, 36, 24] or modifying network structure, for example by using ReLU instead of sigmoid activations. Challenges such as these often prevent new users from training networks to accuracies that approach state of the art and divert attention from designing models that can solve defined problems. Furthermore, backprop is purely sequential in nature – layers must be updated in order, reducing opportunities for parallelization. This limits how well we can exploit the processing power afforded by multi-CPU/GPU setups.

---

\*Asterisk indicates equal contribution from both authors.

This paper seeks to demonstrate that a biologically-motivated algorithm can scale up to the training of large-scale architectures for large databases. Specifically, we will present a procedure that is better suited to parallelization, adjusting synaptic weight parameters with rules that are local in nature (in particular, layers can be updated out of order). The contributions of this work are as follows:

- The algorithm, *recursive local representation alignment* (rec-LRA), is proposed for training large-scale ANNs. Results show that it handles non-differentiable activations, converges faster than backprop, and offers faster training for large-scale benchmarks (ImageNet).
- Strong generalization across several datasets, including the benchmark ImageNet, is demonstrated for architectures trained using rec-LRA. Furthermore, results show that strong performance in a convolutional system can be achieved with far fewer parameters using simple fixed noise maps in place of convolution.

## 2 Related Work

It has long since been a desire of connectionist researchers to develop learning algorithms that simultaneously are biologically-plausible and yield robust generalization to out-of-sample patterns [15, 7, 34, 1]. One key motivation behind the development of alternative algorithms is the removal of the required symmetry between forward pathways for inference and backwards pathways for credit assignment, as is required by backprop. This has also been referred to as the weight-transport problem [10, 22], a strong neuro-biological criticism of backprop as well as one source of its practical issues. Algorithms such as random feedback alignment (FA) [23] and direct feedback alignment (DFA) [25] have shown that learning is possible, surprisingly, even if the feedback pathway is partially decoupled and random, fixed weights are used to transmit derivative signals backward. FA simply replaces the transpose of the feedforward weights in backprop with a similarly-shaped random matrix while DFA directly wire the output layer’s pre-activation derivative to each layer’s post-activation – both algorithms use these random matrices to generate proxies for the partial derivatives normally given by backprop. Under a proposed framework known as discrepancy reduction, it has been shown in [28, 27] that these feedback loops are better suited for generating target representations, entirely removing the global feedback pathway of backprop – a key idea our algorithm builds on. Algorithms such as target propagation [21, 3], which have been shown to fall under the discrepancy reduction framework [27], generate targets through an auto-encoding framework (a decoder attempts to approximate the inverse of a forward encoder’s underlying function).

The idea of local learning, with origins in the classical frameworks of Hebbian [14], anti-Hebbian [8], and competitive learning [32], has slowly begun to gain increased attention in the training of ANNs. Recent proposals have included decoupled neural interfaces [17] and kickback [2]. Furthermore, [41] demonstrated that neural models using simple local Hebbian updates (in a predictive coding framework) could efficiently conduct supervised learning. Earlier approaches that employed local learning included the layer-wise training procedures that were once used to pre-train networks [38, 5, 20, 29]. The problem with these older approaches is that they were greedy—a model was built from the bottom-up, freezing lower-level parameters as higher-level feature detectors were learnt.

## 3 Recursive Local Representation Alignment

In this section, we present our proposed gradient-free learning procedure for training artificial neural networks (ANNs). First, we will define our problem and present notation. Then, we will specify the design of our algorithm in detail.

### 3.1 The Problem & Notation

While our algorithm could be applied to any type of neural architecture (including recurrent ones), in this paper, we will focus on ones that attempt to learn a nonlinear mapping  $f_{\Theta}$  from inputs  $\mathbf{x}$  to outputs  $\mathbf{y}$  (in the case of auto-association  $\mathbf{y} = \mathbf{x}$ ). The input  $\mathbf{x}$  is of dimensionality  $\mathbf{x} \in \mathcal{R}^I$  in the case of simple patterns (as in grey-scale images or text document vectors), where  $I$  is the number of features/observed variables (such as the total number of pixels that compose an image grid or number of unique symbols in a text corpus).<sup>1</sup> In the case of color images, note that  $\mathbf{x} \in \mathcal{R}^{I \times C}$ , where the dimension  $C$  is the number of input channels, e.g., red/blue/green channels. On the other hand,  $\mathbf{y} \in \mathcal{R}^Y$  (a one-hot encoding of any class  $y \in Y$ ), where  $Y$  is the total number of distinct classes/categories in a dataset.

The nonlinear mapping  $f_{\Theta}(\mathbf{x})$  contains a set of learnable parameters housed in the construct  $\Theta$ , which are what any learning procedure, such as backprop, is trying to modify to improve predictive performance. In feedforward networks, a stack of nonlinear transformations, or  $\{f_{\ell}(\mathbf{z}_{\ell-1}; \theta_{\ell})\}_{\ell=1}^L$ , is applied to the input  $\mathbf{x}$ . As an example, if the network is a multilayer perceptron (MLP), each transformation  $\mathbf{z}_{\ell} = f_{\ell}(\mathbf{z}_{\ell-1})$  produces an output  $\mathbf{z}_{\ell}$  from the value  $\mathbf{z}_{\ell-1}$  of the

<sup>1</sup>Vectors and matrices are assumed to be in column-major form.

previous layer with the help of a parameter  $\theta_\ell = \{W_{(\ell-1)\rightarrow\ell}\}$  in the form a weight matrix.  $f_\ell$  is decomposed into two operations (biases omitted for clarity):

$$\mathbf{z}_\ell = \phi_\ell(\mathbf{h}_\ell), \quad \mathbf{h}_\ell = W_{(\ell-1)\rightarrow\ell} \cdot \mathbf{z}_{\ell-1} \quad (1)$$

where  $\phi_\ell$  is an activation function,  $\mathbf{z}_\ell \in \mathcal{R}^H$  is the post-activation of layer  $\ell$  while  $\mathbf{h}_\ell \in \mathcal{R}^H$  is the pre-activation vector of layer  $\ell$ . Note that a matrix multiplication is denoted by  $(\circ \cdot \circ)$ , a Hadamard multiplication is denoted by  $(\circ \otimes \circ)$ , and  $(\circ)^T$  denotes the transpose operator. For convenience, we set  $\mathbf{z}_0 = \mathbf{x}$  (referring to the input vector) and  $\mathbf{z}_L$  is the final output or prediction made by the stacked model  $f_\Theta(\mathbf{x})$ . We have also introduced special notation for our synaptic weight matrices, where  $W_{i\rightarrow j}$  indicates that this parameter matrix connects neurons in layer  $i$  to  $j$ .

For classification/categorization, the output activation is the softmax:  $\mathbf{y} = \phi_L(\mathbf{v}) = \exp(\mathbf{v}) / (\sum_j \exp(\mathbf{v}[j]))$ , where  $j$  indexes scalar elements of a vector. Any element in the output vector, i.e.,  $\mathbf{y}[j] \equiv \phi_L(\mathbf{v})[j] = p(j|\mathbf{v})$ , is the scalar probability of class  $j$ . Generally, the goal of training is to adjust  $\Theta$  to minimize the output loss known as the negative Categorical log likelihood, or  $L(\mathbf{y}, \mathbf{v}) = -\sum_i (\mathbf{y} \otimes \log p(\mathbf{y}|\mathbf{v})) [i]$ .

### 3.2 The Learning Algorithm

The central idea behind our algorithm, recursive local representation alignment (rec-LRA), is that every layer, not just the output layer, has a target and each layer’s parameters/weights are adjusted so that its output moves closer to its target. While this idea is also an aspect of prior work such as target-prop [6, 4, 21], one key difference between rec-LRA and these prior efforts is that rec-LRA chooses targets that are in the “possible representation” of the associated layers. Hence, a layer’s parameters are updated more effectively, i.e., a layer is *not* forced to match a target that is impossible to achieve [27]. Thus, unlike innovations such as target-prop, batch normalization [16], etc, our procedure does *not* need to introduce new layers in the architecture. As a result, it can be viewed either as an alternative to such approaches, or as a complementary technique because it is compatible with other methods, i.e., batch normalization, residual blocks, and any other layer helpful for problem-specific representations that a deep network would need to acquire.

Our algorithm, which builds on and generalizes the ideas in [27, 26] (which were strictly proposed and empirically tested on only fully-connected MLPs), aims to decompose the larger credit assignment problem in neural networks into smaller sub-problems that are not only easier to solve but are also solvable in parallel of each other. As we will show in this section, rec-LRA’s goal is to aggressively decompose the full, underlying directed, acyclic computation/operation graph that defines any stacked neural architecture into small, operation “sub-graphs”. In this paper, we will show how rec-LRA, through its error synaptic network, breaks down a network into its  $L$  individual transformations,  $\{f_\ell(\mathbf{z}_{\ell-1}; \theta_\ell)\}_{\ell=1}^L$ . It follows that this divide-and-conquer behavior of our algorithm naturally facilitates distributed training if high performance computing resources are available.

To specify rec-LRA, we start by defining the function it is ultimately meant to optimize, the *total discrepancy*, which is a “pseudo-energy function” that measures the amount of overall system disorder. Specifically, this function computes the degree of mismatch between the current activity of a neural architecture’s layers and the activity of a set target activities/states. rec-LRA automatically determines the targets but, in principle, the target could come from external sources or be internally generated based on some partially observed external data, representing values that the network’s neuronal processing elements should have taken in order to better predict aspects of its environment.

Under the framework of discrepancy reduction, a neural system is engaged with minimizing the weighted sum of local representational mismatch functions:

$$\mathcal{D}(\Theta) = \sum_{\ell=1}^L \kappa_\ell \mathcal{L}_\ell(\mathbf{y}_\ell, \mathbf{z}_\ell), \quad (2)$$

$$\text{where, } \mathcal{L}_\ell(\mathbf{y}_\ell, \mathbf{z}_\ell) = (\|\mathbf{z}_\ell - \mathbf{y}_\ell\|_p)^q$$

where  $\{\mathbf{y}_1, \dots, \mathbf{y}_\ell, \dots, \mathbf{y}_L\}$  are the layer-wise targets and  $\mathbf{y}_L$  is the output (i.e. it is  $\mathbf{y}$ ). The value  $p$  sets the type of distance function or norm used to compute mismatch between a state’s prediction and the actual target, i.e.,  $p = 2$  is the L2 (Euclidean) norm and  $p = 1$  is the L1 (Manhattan) norm (typically  $q = p$ ). For this study, we set  $p = q = 2$  and choose the Euclidean distance function as our representational mismatch function. The local coefficient  $\kappa_\ell$  is a scalar coefficient that, while typically set to be one for all layers, i.e.,  $\kappa_1 = \dots = \kappa_\ell = \dots = \kappa_L = 1$ , though if set to values less than one, one could simulate different time-scales of parameter evolution within various levels of the model. By taking derivatives of this objective function with respect each layer of neurons, one can then derive vectors of special neurons called “error neurons”, or  $\mathbf{e}_\ell$ . These neurons exclusively measure the difference between a set of post-activity values of one set of neurons  $\mathbf{z}_\ell$  with a corresponding set of target activity values  $\mathbf{y}_\ell$ . These error neurons form the backbone of a learning process conducted in two phases using only forward operations: 1) a target generation

phase aided by the use of synaptic parameters that transmit mismatch signals across the system, and 2) a local weight update step that does not require knowledge of the point-wise derivatives of the ANN’s activation functions.

One particularly powerful and previously unexplored aspect of the discrepancy framework is that the target generation process is independent of the nonlinear mapping  $f_{\Theta}(\mathbf{x})$  we are optimizing. This means that, when conducting credit assignment, we are not constrained by the target network’s forward inference process as is the case for the popular backprop algorithm, which forces us to utilize specific neural circuitry that follows from applying the chain rule of calculus. This ties to backprop’s need for a long global feedback pathway that runs an error signal computed at the output back along the same weights used to forward propagate information [27]. This is not only quite neuro-biologically implausible but it is the central cause of the well-known vanishing/exploding gradient problem [9] since a single error signal must traverse backwards using the same forward weights/parameters along the define the information propagation pathway of  $f_{\Theta}(\mathbf{x})$ , constantly multiplied by the local derivatives of each layer it passes through. In our learning framework, error signals are instead transmitted to the regions/layers of the subgraphs that require them through the use of what we call *skip-error connections*. Skip-error synapses facilitate a direct transmission of mismatch signals computed by neurons at any layer  $i$  directly to any layer  $j$ , serving as a short-circuit pathway. One could also interpret these short-circuit pathways as “error highways”, inspired by the forward synaptic skip connections used to improve the stability of learning deep ANNs via backprop [35].

Given the error neurons and an error synaptic pathway described above, we may now specify how to generate layer-wise targets. In rec-LRA, targets in LRA can be considered to be the latent representations that are more desirable for a network to acquire when trying to learn a useful predictive mapping of  $\mathbf{x}$  to  $\mathbf{y}$ . Generally, in rec-LRA, starting from any error neuron region  $\mathbf{e}_j$ , we compute the target for any immediately connected region/layer  $\mathbf{e}_i$  in the following manner:

$$\mathbf{y}_i = \phi_i(\mathbf{h}_i - \beta \mathbf{d}_i), \text{ where, } \mathbf{d}_i = E_{j \rightarrow i} \cdot \mathbf{e}_j \quad (3)$$

noting that all is required for computing a target at  $i$  is its original pre-activation vector and knowledge of its post-synaptic activation function  $\phi_i(\circ)$ .  $\beta$  is the modulation factor to control the influence of the transmitted error message from node  $j$  to  $i$ . Again, notice that we explicitly indicate the direction of transmission from region  $j$  to  $i$  with the subscript notation  $j \rightarrow i$  for error synapses  $E_{j \rightarrow i}$ . In an MLP,  $\mathbf{h}^i$  could simply be the pre-activation of a layer  $i$  (as in Equation 1) and the post-activity of that layer  $\mathbf{z}^i$  could be computed by applying a non-linear activation function, such as the linear rectifier,  $\phi^i(v) = \max(0, v)$ , or a non-differentiable function such as the signum,  $\phi^i(v) = \text{sign}(v)$ . However,  $\mathbf{h}^i$  could be the output of a complex function, such as a stack of operations, i.e., convolution and max-pooling operators, as in the case of a residual convolutional network.

The full process of computing all of the error neurons embedded in an arbitrary neural architecture is given in Algorithm 1. After running the architecture’s forward pass procedure to gather layer-wise activities, rec-LRA computes mismatch signals by starting at the layer  $L$  and computing the corresponding error neurons  $\mathbf{e}_L$ . From there, rec-LRA retrieves the layer indices of the regions that immediately connect to  $L$  (via an implementation of the function `EXTRACTCHILDRENINDICES( $\circ$ )`), storing these in the array  $\Upsilon$ .  $\Upsilon$  is an un-ordered list of integers, since transmitting the mismatch signal from  $j$  to  $i$  does not depend on the transmission from  $j$  to  $\Upsilon \setminus i$ . This means that the transmission of mismatch signals to each of  $L$ ’s neighbors can be done in parallel if multiple processors are available. For a target region  $i$  connected to  $L$ , rec-LRA will compute its target  $\mathbf{y}_i$  via Equation 3. It will then recursively call itself on that region using the newly computed target, subsequently computing the error neuron vector at  $i$  and further computing




Figure 1: A depiction of recursive-LRA on a feedforward neural, with recursive depth of 3 shown. Green diamonds represent error neurons while blue dashed lines represent error weights that transmit mismatch signals to specific layers in the network.

---

**Algorithm 1** The general recursive local representation alignment algorithm.

---

```

1: Inputs:  $\mathbf{x}, \mathbf{y}, f_{\Theta}(\mathbf{x}), \beta, \gamma$ 
2: // Sub-routine for computing parameter updates for function  $f_{\Theta}(\mathbf{x})$ 
3: function COMPUTEUPDATES( $\mathbf{x}, \mathbf{y}, f_{\Theta}(\mathbf{x})$ )
4:    $\mathcal{H}, \mathcal{Z} = \text{RUNINFERENCE}(\mathbf{x}, \mathbf{y}, f_{\Theta}(\mathbf{x}))$   $\triangleright$  Forward propagate to get pre-activities  $\mathcal{H}$  & post-activities  $\mathcal{Z}$ 
5:    $\mathcal{E}, \vartheta = \text{COMPUTEERRORNEURONS}(\mathbf{y}, f_{\Theta}(\mathbf{x}), \mathcal{H}, \mathcal{Z}), \Delta_{\text{all}} = \emptyset$ 
6:   for  $W_{i \rightarrow j} \in \Theta$  do
7:      $\mathbf{z}_i \leftarrow \mathcal{Z}[i], \mathbf{e}_j \leftarrow \mathcal{E}[j]$ 
8:      $\Delta W_{i \rightarrow j} = \mathbf{e}_j \cdot (\mathbf{z}_i)^T, \Delta_{\text{all}} = \Delta_{\text{all}} \cup \{\Delta W_{i \rightarrow j}\}$ 
9:   for  $E_{j \rightarrow i} \in \Theta$  do
10:     $\mathbf{d}_i \leftarrow \vartheta[i], \mathbf{d}_j \leftarrow \vartheta[j],$ 
11:     $\Delta E_{j \rightarrow i} = -\gamma(\mathbf{e}_i \cdot (\mathbf{d}_j)^T) \Delta_{\text{all}} = \Delta_{\text{all}} \cup \{\Delta E_{j \rightarrow i}\}$ 
12:   return  $\Delta_{\text{all}}$   $\triangleright$  Return full set of parameter updates to  $\Theta$ 
13: // Sub-routine for calculating all error neuron & delta signal vectors for function  $f_{\Theta}(\mathbf{x})$ 
14: function COMPUTEERRORNEURONS( $\mathbf{y}_{\ell}, f_{\Theta}(\mathbf{x}), \mathcal{H}, \mathcal{Z}$ )
15:    $\mathcal{E} = \{\emptyset\} * L, \vartheta = \{\emptyset\} * L$   $\triangleright$  Initialize arrays w/ empty error neurons & delta signals
16:   COMPUTESIGNALS( $L, \mathbf{y}_{\ell}, f_{\Theta}(\mathbf{x}), \mathcal{H}, \mathcal{Z}, \mathcal{E}, \vartheta$ )
17:   return  $\mathcal{E}, \vartheta$ 
18: function COMPUTESIGNALS( $\ell, \mathbf{y}_{\ell}, f_{\Theta}(\mathbf{x}), \mathcal{H}, \mathcal{Z}, \mathcal{E}, \vartheta$ )
19:    $\mathbf{z}_{\ell} = \mathcal{Z}[\ell], \mathbf{e}_{\ell} = (\mathbf{z}_{\ell} - \mathbf{y}_{\ell}), \mathcal{E}[\ell] \leftarrow \mathbf{e}_{\ell}$ 
20:    $\Upsilon \leftarrow \text{EXTRACTCHILDRENINDICES}(\ell, f_{\Theta}(\mathbf{x}))$   $\triangleright$  Extract list of children node indices for error node  $\ell$ 
21:   // Recursive Case: Traverse into each children error node & compute its error vector
22:   if  $\Upsilon$  is not  $\emptyset$  then
23:     for  $i \in \Upsilon$  do
24:        $\mathbf{h}_i = \mathcal{H}[i], \mathbf{d}_i = E_{\ell \rightarrow i} \cdot \mathbf{e}_{\ell}, \mathbf{y}_i = \phi_i(\mathbf{h}_i - \beta \mathbf{d}_i), \vartheta[i] \leftarrow \mathbf{d}_i$ 
25:       COMPUTESIGNALS( $i, \mathbf{y}_i, \mathcal{G}_{\Theta}, f_{\Theta}(\mathbf{x}), \mathcal{H}, \mathcal{Z}, \mathcal{E}, \vartheta$ )
26:   // Base Case: No children error neurons, so no need to further update graph at node  $\ell$ 

```

---

targets for any regions  $\Upsilon$  connected to  $i$  and so on and so forth. The base case for termination in full rec-LRA is simply the situation when  $\Upsilon = \emptyset$ , i.e., there are no regions that immediately connect to  $i$ . Once all error neuron vectors have been computed, we can calculate updates to all parameters of not only the neural architecture but also to each error matrix used to transmit  $\gamma$  is a decay factor (typically set close to 1.0) meant to ensure that the error weights change more slowly than the forward weights.

While the pseudocode in Algorithm 1 first computes the error neurons ( $\text{COMPUTEERRORNEURONS}(\circ)$ ) then calculates parameters updates (lines 6-11 in  $\text{COMPUTEUPDATES}(\circ)$ ) after, one could actually merge the two functions together and immediately compute the updates for any incoming model weights  $W_{k \rightarrow i}$  that connect to region  $i$  (as well any relevant error synapses  $E_{j \rightarrow i}$ ). Furthermore, even though the algorithm as presented would execute each recursive call sequentially (in the sub-routine  $\text{COMPUTESIGNALS}(\circ)$ , given that the transmission of error from  $j$  to  $i$  is independent of that from  $j$  to  $\Upsilon \setminus i$ , one could allocate each recursive to a cluster/set of CPUs/GPUs that are dedicated with generating targets for all parts of the operator graph that the recursive call will see.<sup>2</sup> This design highlights one of rec-LRA's key strengths – it compute targets and parameter updates in a divide-and-conquer approach utilizing the transmission pathways defined by error connectivity.

### 3.3 Residual Neural Networks and rec-LRA

While rec-LRA could be applied to any neural architecture, in this paper consider the case of the residual (neural) network (ResNet) [12, 13]. Residual networks, which have recently achieved state-of-the-art performance on several popular vision benchmarks, are architectures that are composed of many hidden layers wired together with a special forward connectivity pattern. Specifically, residual networks utilize skip/shortcut connections that allow the forward propagation of information to potentially jump over some hidden layers, specifically those that might not prove useful

---

<sup>2</sup>One important assumption made by Algorithm 1 is that no cycles have been introduced when specifying the error transmission pathways, otherwise, extra steps would need to be taken to ensure that the independence of error transmission pathways is maintained after a user specifies the set of error transmission matrices.




Figure 2: Rec-LRA applied to 2 blocks of a ResNet. Left (of blue line) depicts target creation & right depicts update calculation. in mapping  $x$  to  $y$ . Formally, the layers in the network that permit a residual mapping are defined as:

$$z_\ell = f_\ell(z_{\ell-1}; \theta_\ell) + z_{\ell-g}$$

where  $g$  controls the length of the gap/skip, typically of size 2 or 3. The idea behind the formulation above is that, in the event that directly fitting the transformation function  $f_\ell(z_{\ell-1}; \theta_\ell)$  is too challenging, the residual mapping (as indicated by the second term of the equation) will be easier to optimize. This is, in effect, gives the network the choice of retaining the input if it finds that a particular layer or layers are not needed. The transformation  $f_\ell(z_{\ell-1}; \theta_\ell)$  could be as simple as a linear transformation or a stack of fully-connected layers (as in Equation 1). In computer vision, it is often formulated as a residual “block”, where a stack of multiple operations is housed, e.g., convolutions, the relu activation function ( $\mathbf{v} = \max(0, \mathbf{v})$ ), pooling, normalization layers, etc. A typical residual block is depicted in Figure 3.




Figure 3: Residual block.

Training a residual network with rec-LRA is rather straightforward – we may exploit the block-based structure of the architecture in crafting the error message transmission pathways. If, for example, a residual block is a stack of nonlinear transformations (or a 3-hidden layer MLP as depicted in Figure 2), we can choose to embed a vector of error neurons at the output of each residual block and wire them to the output error neurons at layer  $L$ . In the case of the two residual blocks depicted in Figure 2, we would use skip-error connections  $E_{L \rightarrow i}$  and  $E_{L \rightarrow j}$ . Wiring skip-error connections in this way means that, under rec-LRA, we are treating each residual block as a computational subgraph (which maps a representation  $z_{i-g}$  to  $z_i$ ). In this case, once a skip-error connections wired to each a block generates its desired target, rec-LRA will recursively enter the block to compute its internal error neurons and weight updates, independently of the blocks above and below, effectively decoupling its update calculation from the rest of the blocks. In treating the residual blocks as decoupled computation graphs, one could view the output of each block as a “meta-representation” (in Figure 2, these would be  $z_i, z_j$ ), or a particular post-activation layer that serves as the focus of LRA’s target generation process while the other layers within it would be could simply be viewed as computational “support” layers, i.e.,  $\{z_{i-1}, z_{i-2}\}$  and  $\{z_{j-1}, z_{j-2}\}$  in Figure 2. In Algorithm 2, we formally depict how rec-LRA would compute targets and weight updates to a residual network (composed of fully connected layers) with a skip  $g$ .

While rec-LRA could continue to decompose each block using additional skip-error connections and compute weight updates using the simple local rule presented in Algorithm 2, one could opt to use a different learning algorithm for the internal operations of the residual blocks. Specifically, one could use rec-LRA to generate meta-representation targets for the residual blocks, e.g.,  $z_i$ , and then employ another procedure such as backprop (treating the block’s output error neuron vector as a proxy for  $\frac{\partial \mathcal{L}}{\partial z_i}$ ) or a local Hebbian rule [14] to compute the actual weight adjustments.

**Replacing Convolution with Fixed Perturbation:** To further save on computation, we replaced the convolutional operator with a fixed noise “pseudo-convolution”, which was proposed in [18] (referred to as a “perturbative layer”). As

---

**Algorithm 2** Rec-LRA (depth 2) applied to L-layered network  $f_{\Theta}(\mathbf{x})$  with residual gap  $g$ .

---

```

1: Inputs:  $\mathbf{x}, \mathbf{y}, \Theta = \{W_1, \dots, W_L\}, \beta, \gamma, g$ 
2:    $\Theta_E = \{E_{L \rightarrow (L-1)}, \dots, E_{i \rightarrow j}, \dots, E_{L \rightarrow 1}\}$ 
3: // Inference procedure for network  $f_{\Theta}(\mathbf{x})$ 
4: function RUNMODEL( $\mathbf{x}, \Theta$ )
5:   for  $\ell = 1$  to  $L$  do
6:     if  $\ell \bmod g \equiv 0$  then
7:        $\mathbf{h}_{\ell} = W_{\ell} \cdot \mathbf{h}_{\ell-1} + \mathbf{z}_{\ell-1}$ 
8:     else
9:        $\mathbf{h}_{\ell} = W_{\ell} \cdot \mathbf{h}_{\ell-1}$ 
10:     $\mathbf{z}_{\ell} = \phi_{\ell}(\mathbf{h}_{\ell})$ 
11:   return  $\mathcal{Z} = \{\mathbf{z}_0, \mathbf{z}_1, \dots, \mathbf{z}_L\}$ 
12: // Computes error neurons given activities of  $f_{\Theta}(\mathbf{x})$ 
13: function COMPUTEERRORNEURONS( $\mathbf{y}, \Theta, \Theta_E, \mathcal{Z}$ )
14:    $\mathbf{y}_L = \mathbf{y}, \mathbf{e}_L = \mathbf{z}_L - \mathbf{y}_L$ 
15:   for  $\ell = (L - 1)$  to  $1$  do
16:     if  $\ell \bmod g \equiv 0$  then
17:        $\mathbf{d}_{\ell} = E_{L \rightarrow \ell} \cdot \mathbf{e}_L$ 
18:     else
19:        $\mathbf{d}_{\ell} = E_{(\ell+1) \rightarrow \ell} \cdot \mathbf{e}_{\ell+1}$ 
20:      $\mathbf{y}_{\ell} = \phi_{\ell}(\mathbf{h}_{\ell} - \beta \mathbf{d}_{\ell}), \mathbf{e}_{\ell} = \mathbf{z}_{\ell} - \mathbf{y}_{\ell}$ 
21:   return  $\mathcal{E} = \{\mathbf{e}_1, \dots, \mathbf{e}_L\}, \Upsilon = \{\mathbf{d}_1, \dots, \mathbf{d}_L\}$ 
22: // Computes updates given error neurons of  $f_{\Theta}(\mathbf{x})$ 
23: function COMPUTEUPDATES( $\mathcal{E}, \mathcal{Z}, \Upsilon$ )
24:   for  $\ell = 1$  to  $L$  do
25:      $\Delta W_{\ell} = \mathbf{e}_{\ell} \cdot (\mathbf{z}_{\ell-1})^T$ 
26:     if  $\ell > 1$  then
27:       if  $\ell \bmod g \equiv 0$  then
28:          $\Delta E_{\ell} = -\gamma(\mathbf{d}_{\ell-1} \cdot (\mathbf{e}_L)^T)$ 
29:       else
30:          $\Delta E_{\ell} = -\gamma(\mathbf{d}_{\ell-1} \cdot (\mathbf{e}_{\ell})^T)$ 
31:   return  $\{\Delta W_1, \Delta W_2, \Delta E_2, \dots, \Delta W_L, \Delta E_L\}$ 

```

---

was shown in [18], the pseudo-convolution is not only drastically faster than actual convolutional but the generalization performance of the underlying model using it is comparable to one with convolution. A pseudo-convolution is computed as follows:

$$\mathbf{z}_{\ell}^c = \sum_{m=1}^M w_{\ell}^m \phi_r(\mathbf{z}_{\ell-1}^m + \mathbf{n}_{\ell-1}^m),$$

where,  $\phi_r(v) = \max(0, v)$

where  $w_m^{\ell}$  is a scalar weight that is applied to its corresponding noise map. In the above formula, we see that  $M$  noise maps  $\mathbf{n}_{\ell}^m$  must be cycled through in order to compute the final desired output channel (map)  $\mathbf{z}_{\ell}^c$ . The idea is that, for the price of the memory required to store the pre-generated  $M$  noise maps (the elements of each are each sampled from a centered Gaussian distribution,  $\sim \mathcal{N}(0, \sigma^2)$ ), we may side-step the need for learn-able kernel parameters which is required for a convolution operation (per filter that is applied to an incoming image channel). The only parameters in a pseudo-convolution that require updating are the linear combination weights (one update per scalar weight applied to each noise map). Under rec-LRA, which would embed error neurons right next to  $\mathbf{z}_{\ell}^c$ , the update for the  $m$ th noise map weight  $w_m^{\ell}$  would be:

$$\Delta w_{\ell}^m = \sum_i \sum_j \left( \mathbf{e}_{\ell} * (\phi_r(\mathbf{z}_{\ell-1}^m + \mathbf{n}_{\ell-1}^m))^T \right) [i, j]$$

where we collapse the update by summing along all dimensions in order to obtain a scalar update for the weight  $w_{\ell}^m$ .

Table 1: Generalization error of networks trained on the MNIST, Fashion MNIST (FMNIST), CIFAR-10 benchmarks.

Algorithm	MNIST		FMNIST		CIFAR	
	Train	Test	Train	Test	Train	Test
Backprop	<b>0.00%</b>	<b>1.48%</b>	<b>12.10%</b>	<b>12.98%</b>	CNN, Backprop	1.39 31.87
TP	<b>0.00%</b>	<b>1.86%</b>	21.078%	19.66%	TP	28.69 39.47
E-Prop	7.59%	9.21%	16.56%	20.97%	FA	17.46 37.44
LRA-E	0.16%	1.97%	<b>9.84%</b>	<b>12.31%</b>	DFA	32.74 44.41
FA	<b>0.00%</b>	<b>1.85%</b>	<b>12.09%</b>	<b>12.89%</b>	Backprop (impl)	<b>5.00</b> <b>5.94</b>
DFA	0.85%	2.75%	12.58%	13.09%	rec-LRA (ours)	<b>5.88</b> <b>6.12</b>
rec-LRA, tanh (ours)	<b>0.00%</b>	<b>1.82%</b>	<b>6.57%</b>	<b>11.87%</b>	ImageNet	
rec-LRA, lrelu (ours)	0.22%	2.26%	8.95%	14.13%	Algorithm	Top-1 Top-5
rec-LRA, elu (ours)	0.09%	1.93%	9.39%	13.17%	TP	98.34 94.56
rec-LRA, sign (ours)	0.85%	2.33%	12.42%	14.879%	FA	93.08 82.54
					Backprop (impl)	28.15 9.81
					rec-LRA (ours)	30.48 11.97




Figure 4: Error for networks on MNIST (top) &amp; Fashion MNIST (middle). Top-1 error reported for ImageNet ResNets (bottom).

## 4 Experiments

In this section, we experiment with our proposed algorithm adapted to feedforward networks and compare it to results reported for other backprop-alternatives. Specifically, we experimented with rec-LRA adapted to fully-connected MLPs and residual convolutional network structures.

For optimization, we employed the Adam [19] adaptive learning rate, using a learning rate of  $\lambda = 2e - 4$  (tuned using validation performance for each dataset). For the MNIST and Fashion MNIST datasets, updates to parameters, whether they were calculated via backprop or rec-LRA, were estimated over mini-batches of size 32 and layers were initialized according to an element-wise, zero-mean Gaussian distribution with standard deviation  $\sigma$  selected in the range of  $[0.025 - 0.1]$  (tuned using validation performance). Models with 5 or 8 layers of 256 neurons were trained over 500 epochs.<sup>3</sup> Models with 3 layers of 1024 neurons were trained for 100 epochs.

For the CIFAR-10 and ImageNet benchmarks, parameter updates were estimated with mini-batches of 10 samples and Adam was used with learning rate of  $\lambda = 1e - 4$ . We used additive noise set to a level of 0.1, 256 perturbation masks per layer for ImageNet, and 160 perturbation masks per layer for CIFAR-10. We trained each model for 100 epochs and tuned individual meta-parameters based on validation performance. The rest of our configuration settings were set to be similar to [18] in order to ensure a fair comparison among models. Layers were initialized with a unit Gaussian distribution<sup>4</sup> For all architectures and algorithms in all experiments, we re-projected the weight updates back to a Gaussian ball of radius  $c$  (as in [30]). Further details can be found in the appendix.

**MNIST & Fashion MNIST:** This dataset<sup>5</sup> contains  $28 \times 28$  images with gray-scale pixel feature values in the range of  $[0, 255]$ . The only preprocessing applied to this data is to normalize the pixel values to the range of  $[0, 1]$  by dividing them by 255. On the other hand, Fashion MNIST (FMNIST) [39] was proposed as a challenging drop-in replacement for MNIST. Fashion MNIST (pre-processed the same as MNIST) contains  $28 \times 28$  grey-scale images depicting clothing

<sup>3</sup>We chose this configuration and # of epochs to ensure we could compare to related prior work [3].

<sup>4</sup>Xavier and orthogonal initialization schemes were found to yield unsatisfactory performance.

<sup>5</sup>Available at the URL: <http://yann.lecun.com/exdb/mnist/>.



items, each image associated to one of 10 item classes. Training had 60000 samples, testing had 10000, and validation of 2000 samples was drawn from the training set.

In Table 1, we report our classification error on both training and test sets for rec-LRA and compare to prior reported results. Prior results have been reported for backprop (Backprop) as well as relevant biologically-motivated, gradient-free algorithms such as feedback alignment (FA), direct feedback alignment (DFA), a error-driven local representation alignment (LRA-E), equilibrium propagation (E-Prop), and target propagation (TP) (results taken from [3] and [26]). For the rec-LRA results, we report 4 variations, each using a different activation function. To be comparable to prior work, the first variant of rec-LRA utilizes hyperbolic tangent units (rec-LRA, tanh). The bottom three variants entailed using leaky rectifier units (rec-LRA, lrelu), exponential linear units (rec-LRA, elu), and finally, signum units (rec-LRA, sign) in order to investigate rec-LRA’s effectiveness in training networks with non-differentiable functions. We observe that rec-LRA outperforms all of the other algorithms on Fashion MNIST, including backprop. On MNIST, rec-LRA outperforms all of the other gradient-free alternatives but does not beat out backprop. While the signum networks do not reach the performance of the topmost networks, they are not among the worst performing, which offers encouraging evidence that non-differentiable networks can make viable classifiers.

However, we further analyzed the training dynamics of more complex, nonlinear networks composed of 8 layers of either 256 logistic sigmoid or tanh neurons trained via backprop and rec-LRA over a period of 100 epochs. Deep sigmoidal units are known to be incredibly difficult to train, giving rise to the well-known vanishing gradient problem [9], especially if naive Gaussian initializations are used. In Figure 4, on both MNIST and Fashion MNIST, we observe that rec-LRA is not only able to successfully train networks of both kinds of units with the same initialization but is also able to converge sooner. The fact that this result holds for the networks with tanh units, which are friendlier to a backprop-centric optimization, offers some evidence of rec-LRA’s potential robustness and stability.

**CIFAR-10:** The CIFAR-10 dataset contains 50,000 training and 10,000 test images, labeled as one of 10 possible image categories. The images are of size  $32 \times 32$  pixels. 5,000 of the training images were set aside to measure validation loss and accuracy. In terms of pre-processing, global contrast normalization was first applied to image, where each color channel’s pixel mean was subtracted from itself. ZCA whitening was then subsequently applied, where: 1) the image data was centered and rotated onto its principle components, 2) the principle components were normalized, and 3) the image data was finally rotated back.

While this dataset is far more challenging than that of MNIST, we observe in Table 1 that rec-LRA outperforms networks trained with other gradient-free methods, i.e, target prop and feedback alignment. Furthermore, rec-LRA comes quite close to the performance the same architecture trained with backprop, offering evidence of its ability to handle a challenging problem involving color images. In the appendix, we further dissect the networks’ predictions and visualize the top-most latent representations acquired by each network via t-SNE [37].

**ImageNet:** The large-scale benchmark ImageNet [33], specifically the ILSVRC-2010 subset, is contains over 1.2 million images, of size  $224 \times 224$ , each contain one out of 1000 different categories. Given that the number of classes is large, it is convention to report two types of error rates: top-1 and top-5. The top-5 error rate is the fraction of test images for which the correct label is not among the 5 classes considered most probable by the model being evaluated. In Table 1, we observe that the network trained via rec-LRA not only outperformed networks trained via gradient-free methods, but again, comes quite close to the performance of the backprop-trained architecture in terms of both top-1 and top-5 test error.

Furthermore, we measured wall-clock training time for both both networks to determine if rec-LRA training offered a speed-up (since it does not require traversing down a long global feedback pathway nor does it require computing activation function derivatives). Notably, in terms of total training run-time over 90 epochs using a small set of 8 V100 GPUs, the backprop ResNet took 3 hours and 45 minutes (min) to train (speed was about 2.5-2.7 min/epoch) while rec-LRA took 2.127 min/epoch, training over the course of 3 hours and 12 min. In Figure 4 (bottom), we see that rec-LRA does reach lower validation error sooner than backprop (though this result is not as obvious as it was for MNIST/Fashion MNIST). Furthermore, rec-LRA converges more smoothly than the backprop-trained ResNet. State-of-the-art performance of deep networks on ImageNet is much lower [40] than what has been reached with gradient-free algorithms such as our own and in [3]. However, our aim was to demonstrate that a gradient-free algorithm can indeed yield strong generalization on difficult, large-scale datasets. Effort to utilize the various heuristics/tweaks that have been applied to carefully tune modern-day networks to ImageNet over the years would certainly further boost our own ResNet’s performance, potentially approaching state-of-the-art error.

## 5 Conclusions

In this paper, we proposed a gradient-free learning algorithm, recursive local representation alignment algorithm (rec-LRA), for training deep neural architectures. rec-LRA generalizes as well as backprop and outperforms other current gradient-free procedures across several datasets, notably on the massive-scale benchmark ImageNet. Furthermore, it offers improved convergence due to faster, parallelizable weight updates, as shown in our experiments. As a result, this work offers empirical evidence that a backprop-free procedure can indeed scale up to larger datasets.

## References

- [1] ALIAS PARTH GOYAL, A. G., KE, N., GANGULI, S., AND BENGIO, Y. Variational walkback: Learning a transition operator as a stochastic recurrent net. In *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, Eds. Curran Associates, Inc., 2017, pp. 4392–4402.
- [2] BALDUZZI, D., VANCHINATHAN, H., AND BUHMANN, J. M. Kickback cuts backprop’s red-tape: Biologically plausible credit assignment in neural networks. In *AAAI (2015)*, pp. 485–491.
- [3] BARTUNOV, S., SANTORO, A., RICHARDS, B., MARRIS, L., HINTON, G. E., AND LILLICRAP, T. Assessing the scalability of biologically-motivated deep learning algorithms and architectures. In *Advances in Neural Information Processing Systems (2018)*, pp. 9390–9400.
- [4] BENGIO, Y. How auto-encoders could provide credit assignment in deep networks via target propagation. *CoRR abs/1407.7906* (2014).
- [5] BENGIO, Y., LAMBLIN, P., POPOVICI, D., LAROCHELLE, H., ET AL. Greedy layer-wise training of deep networks. *Advances in neural information processing systems 19* (2007), 153.
- [6] CARREIRA-PERPIÑÁN, M. Á., AND WANG, W. Distributed optimization of deeply nested systems. *CoRR abs/1212.5921* (2012).
- [7] CHALASANI, R., AND PRINCIPE, J. C. Deep predictive coding networks. *arXiv preprint arXiv:1301.3541* (2013).
- [8] FÖLDIAK, P. Forming sparse representations by local anti-hebbian learning. *Biological cybernetics 64*, 2 (1990), 165–170.
- [9] GLOROT, X., AND BENGIO, Y. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* (2010), pp. 249–256.
- [10] GROSSBERG, S. Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science 11*, 1 (1987), 23 – 63.
- [11] HE, K., ZHANG, X., REN, S., AND SUN, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the 2015 IEEE International Conference on Computer Vision (ICCV)* (2015).
- [12] HE, K., ZHANG, X., REN, S., AND SUN, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition* (2016), pp. 770–778.
- [13] HE, K., ZHANG, X., REN, S., AND SUN, J. Identity mappings in deep residual networks. In *European conference on computer vision* (2016), Springer, pp. 630–645.
- [14] HEBB, D. O. The organization of behavior; a neuropsychological theory. *A Wiley Book in Clinical Psychology*. (1949), 62–78.
- [15] HINTON, G. E. Training products of experts by minimizing contrastive divergence. *Neural computation 14*, 8 (2002), 1771–1800.
- [16] IOFFE, S., AND SZEGEDY, C. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning* (2015), pp. 448–456.
- [17] JADERBERG, M., CZARNECKI, W. M., OSINDERO, S., VINYALS, O., GRAVES, A., AND KAVUKCUOGLU, K. Decoupled neural interfaces using synthetic gradients. *arXiv preprint arXiv:1608.05343* (2016).
- [18] JUEFEI-XU, F., NARESH BODDETI, V., AND SAVVIDES, M. Perturbative neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2018), pp. 3310–3318.

- [19] KINGMA, D. P., AND BA, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [20] LEE, C.-Y., XIE, S., GALLAGHER, P., ZHANG, Z., AND TU, Z. Deeply-Supervised Nets. *arXiv:1409.5185 [cs, stat]* (2014).
- [21] LEE, D.-H., ZHANG, S., FISCHER, A., AND BENGIO, Y. Difference target propagation. In *Proceedings of the 2015th European Conference on Machine Learning and Knowledge Discovery in Databases - Volume Part I* (Switzerland, 2015), ECMLPKDD'15, Springer, pp. 498–515.
- [22] LIAO, Q., LEIBO, J. Z., AND POGGIO, T. A. How important is weight symmetry in backpropagation? In *AAAI* (2016), pp. 1837–1844.
- [23] LILLICRAP, T. P., COWNDEN, D., TWEED, D. B., AND AKERMAN, C. J. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications* 7 (2016), 13276.
- [24] MISHKIN, D., AND MATAS, J. All you need is a good init. *CoRR abs/1511.06422* (2015).
- [25] NØKLAND, A. Direct feedback alignment provides learning in deep neural networks. In *Advances in Neural Information Processing Systems* (2016), pp. 1037–1045.
- [26] ORORBIA, A. G., AND MALI, A. Biologically motivated algorithms for propagating local target representations. In *Proceedings of the AAAI Conference on Artificial Intelligence* (2019), vol. 33, pp. 4651–4658.
- [27] ORORBIA, A. G., MALI, A., KIFER, D., AND GILES, C. L. Deep credit assignment by aligning local representations. *arXiv preprint arXiv:1803.01834* (2018).
- [28] ORORBIA II, A. G., HAFFNER, P., REITTER, D., AND GILES, C. L. Learning to adapt by minimizing discrepancy. *arXiv preprint arXiv:1711.11542* (2017).
- [29] ORORBIA II, A. G., REITTER, D., WU, J., AND GILES, C. L. Online learning of deep hybrid architectures for semi-supervised categorization. In *Machine Learning and Knowledge Discovery in Databases (Proceedings, ECML PKDD 2015)*, vol. 9284 of *Lecture Notes in Computer Science*. Springer, Porto, Portugal, 2015, pp. 516–532.
- [30] PASCANU, R., MIKOLOV, T., AND BENGIO, Y. On the difficulty of training recurrent neural networks. In *International Conference on Machine Learning* (2013), pp. 1310–1318.
- [31] RUMELHART, D. E., HINTON, G. E., AND WILLIAMS, R. J. Learning representations by back-propagating errors. *nature* 323, 6088 (1986), 533–536.
- [32] RUMELHART, D. E., AND ZIPSER, D. Feature discovery by competitive learning. *Cognitive science* 9, 1 (1985), 75–112.
- [33] RUSSAKOVSKY, O., DENG, J., SU, H., KRAUSE, J., SATHEESH, S., MA, S., HUANG, Z., KARPATY, A., KHOSLA, A., BERNSTEIN, M., ET AL. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115, 3 (2015), 211–252.
- [34] SCELLIER, B., AND BENGIO, Y. Equilibrium propagation: Bridging the gap between energy-based models and backpropagation. *Frontiers in computational neuroscience* 11 (2017), 24.
- [35] SRIVASTAVA, R. K., GREFF, K., AND SCHMIDHUBER, J. Highway networks. *arXiv preprint arXiv:1505.00387* (2015).
- [36] SUSSILLO, D. Random walks: Training very deep nonlinear feed-forward networks with smart initialization. *CoRR abs/1412.6558* (2014).
- [37] VAN DER MAATEN, L. Barnes-hut-sne. *arXiv preprint arXiv:1301.3342* (2013).
- [38] VINCENT, P., LAROCHELLE, H., BENGIO, Y., AND MANZAGOL, P.-A. Extracting and composing robust features with denoising autoencoders. In *Proceedings of the 25th international conference on Machine learning* (2008), ACM, pp. 1096–1103.
- [39] XIAO, H., RASUL, K., AND VOLLGRAF, R. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747* (2017).
- [40] XIE, Q., HOVY, E., LUONG, M.-T., AND LE, Q. V. Self-training with noisy student improves imagenet classification. *arXiv preprint arXiv:1911.04252* (2019).
- [41] XIE, X., AND SEUNG, H. S. Equivalence of backpropagation and contrastive hebbian learning in a layered network. *Neural computation* 15, 2 (2003), 441–454.

## Appendix

### Additional Analysis of Model Predictions:

In Table 2, we dissect the networks’ predictions (beyond accuracy) by analyzing their confusion matrices on the test set – we calculate precision (Prec), recall (Rec), and the F1 score (the harmonic mean between Prec and Rec). In terms of these metrics, it appears that rec-LRA is a bit weaker in recall compared to backprop, though its precision is quite close to that of backprop. We speculate that the small gap in performance could be closed with a more rigorous tuning of the meta-parameters of rec-LRA on the validation set. Nonetheless, rec-LRA’s strong generalization on CIFAR-10 already offers evidence of its ability to scale up to a more challenging problem involving color images.

Table 2: Generalization performance measurements of networks trained on CIFAR-10.

Algorithm	Acc	F1	Prec	Rec
Backprop (impl)	91.15%	90.24	89.51	90.58
rec-LRA (ours)	90.88%	89.96	89.20	88.61

Table 3: Generalization analysis of networks trained on ImageNet.

Algorithm	Accuracy	F1	Precision	Recall
Backprop	71.84%	70.04	70.12	69.96
rec-LRA	69.52%	67.82	68.12	67.52

As was done for CIFAR-10, in Table 3, we further analyze the ImageNet networks’ prediction performance. This time we observe a bit of a larger gap, especially in terms of recall. However, with minimal tuning of the ResNet trained via rec-LRA, its generalization performance is quite impressive. We hypothesize more rigorous/careful tuning will improve performance across all metrics.

### t-SNE Visualization of CIFAR-10 Models

We visualize, the top-most latent representations acquired by the network trained by backprop against rec-LRA via t-SNE [37] in Figure 5. Qualitatively, we observe that rec-LRA does indeed learn a good separation/clustering of classes in its latent representations (just as backprop does).




Figure 5: t-SNE viz. of Resnet – backprop (left), rec-LRA (right).

### Update Re-Projection

For all architectures and algorithms in all experiments of this paper, we re-projected the weight updates back to a Gaussian ball of radius  $c$ . Specifically, this normalization proceeds formally as follows:

$$Nm(\Delta, c) = \left\{ \frac{c}{\|\Delta\|} \Delta, \text{ if } \|\Delta\| \geq c, \text{ and } \Delta, \text{ if } \|\Delta\| < c \right\}$$

where  $\Delta$  is any parameter update matrix returned by a learning algorithm. We found that gradient re-projection was useful in ensuring stable and consistent training.